Characterizing x86 processors for industry-standard servers: AMD Opteron and Intel Xeon

technology brief, 2nd Edition

Abstract

The HP ProLiant server portfolio includes systems using the Intel® Xeon™ family of x86 processors and systems using AMD Opteron™ x86 processors. To help customers understand the implications of these processor architectures on the system platform, this paper compares and contrasts the Xeon processors with the Opteron processor. It discusses both the 32-bit and 64-bit computing space for those x86 processors.

This paper is intended for IT professionals familiar with industry-standard server technology.

Acronyms in text

The following acronyms are used in the text of this document.

Acronym abbreviation	Term
2P	Containing two microprocessors
4P	Containing four microprocessors
8P	Containing eight microprocessors
AGU	Address generation unit
ALU	Arithmetic logic unit
CPU	Central processing unit, or microprocessor
DDR	Double data rate
DIMM	Dual inline memory module
DP	Dual processor
DRAM	Dynamic random access memory
FSB	Front-side bus
Gb	Gigabit
GB	Gigabyte
GPR	General purpose register
KB	Kilobyte
I/O	Input/output – commonly referring to devices such as keyboards mice, video graphics, etc.
IT	Information technology
MHz	Megahertz
MP	Multiprocessor
nm	Nanometer
ns	Nanosecond
OS	Operating system
PCI	Peripheral component interconnect, a standard for connecting peripherals to a computer
SDRAM	Synchronous dynamic random access memory
USB	Universal serial bus

Introduction

Since the early 1980's, industry-standard computers have used the x86 instruction set architecture. During the last 20 years, the x86 architecture has been expanded with more instructions and additional registers to enable easier floating point and multimedia calculations. Throughout these expansions, the x86 architecture has maintained a high level of compatibility with 16-bit and 32-bit software, providing important backward compatibility with software applications. Furthermore, the x86 architecture has offered continually increasing performance at decreasing cost levels.

Intel processor variations

There are many versions of Intel x86 processors, configured for different markets. As of this publication, the most recent models of the Intel x86 processors for industry-standard servers are the Pentium® 4, Xeon™, and the Xeon processor MP using 90 nm technology. All the Intel processors listed in Table 1 use the NetBurst® Architecture and support Hyper-Threading Technology, but they have varying core frequencies, system bus frequencies, and amounts of cache. Hyper-Threading Technology[1] provides multi-thread level parallelism in a single processor core. The processors support Extended Memory 64-bit technology (EM64T) that allows 64-bit operating systems and applications to run natively. The use of 64-bit extension technology will enable IT organizations to deploy common platforms for both 32-bit and 64-bit computing, and to move to 64-bit computing gradually as it benefits their businesses.

Table 1. Intel processor specifications (maximum values are given as October 2005)

Processor	Core frequency (GHz)	L1 cache (KB)	L2 cache (MB)	L3 cache (MB)	Front side bus speed (MHz)	Feature size (nm)	64-bit support
Pentium 4 Supporting Hyper-Threading technology	3.8	16 K data 12K µop instruction	2	None	800	90	Yes
Xeon	3.8	16 K data 12K µop instruction	2	None	800	90	Yes
Xeon MP (code-named Cranford)	3.66	8 K data 12K µop instruction	1	none	667	90	Yes
Xeon MP (code-named Potomac)	3.3	8 K data 12K µop instruction	1	8	667	90	Yes

Intel has introduced dual-core technology in certain Pentium 4 processors designed for the desktop market. Intel has announced dual-core processors for the server market in the fourth quarter of 2005. These first dual-core processors are expected to use the existing NetBurst/HyperThreading architecture without fundamental changes to the micro-architecture.

A dual-core processor is a single physical package that contains two full processor execution cores on one or more die. The physical implementation of that processor will vary by processor vendor and also may vary over time. Dual-core architecture delivers higher performance for applications that are

[1] Additional information about Hyper-Threading is available in the technology brief titled "The Intel® processor roadmap for industry-standard servers," at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00164255/c00164255.pdf

multi-threaded, because the compute threads can be split between the two cores, leveraging the benefits of multi-threaded applications and reducing the problem of multiple threads competing for the same execution resources. Having two processor cores in a single package improves performance significantly while maintaining a similar power budget to single-core processors.

Customers should be aware of possible changes in software licensing when using dual-core processors. As of this writing, major OS vendors such as Microsoft[2] are treating dual-core processors as performance improvements to a single processor and not as two separate processors, depending on OS version. Customers should check with their OS and application vendors to determine their particular licensing requirements.

AMD processor variations

In 2003, AMD introduced its eighth-generation x86 processor, the Opteron. The Opteron was the first processor in the industry to deliver 64-bit extensions (the AMD64 instruction set) for running 64-bit operating systems and applications. In 2005, AMD was the first vendor to introduce dual-core technology for the x86 market with the Opteron dual-core processor version. The Opteron processor uses a consistent architecture—the Direct Connect Architecture—whether the processor is single-core or dual-core or is used in 2P or 4P servers. As shown in Table 2, this architecture includes:

- 1 MB L2 cache, 64 KB of L1 instruction cache, and 64 KB of L1 data cache per core
- 64-bit architecture that can use either 32-bit (x86) or 64-bit (AMD64) instruction sets
- Three HyperTransport interfaces for fast, point-to-point connections between processors and to I/O
- Integrated 128-bit memory controller that runs at the speed of the processor and supports DDR memory up to PC3200

Because the memory controller is integrated and the Opteron processors use the HyperTransport (point-to-point) architecture, there is no need for a front-side bus system such as the Xeon uses.

Table 2. AMD processor specifications (maximum values as of October 2005)

Processor	Core frequency (GHz)	L1 cache (KB)	L2 cache (MB)	L3 cache (MB)	Front side bus speed	Feature size (nm)	64-bit support
Opteron 2YY (single core—for up to 2Pservers)	2.8	64K data 64 K instruction	1024	None	N/A - Integrated memory controller	90	Yes
Opteron 2YY (dual-core—for up to 2P servers)	2.4	64K data 64 K instruction	1024	None	N/A - Integrated memory controller	90	Yes
Opteron 8YY (single core—for up to 8P servers)	2.8	64K data 64 K instruction	1024	None	N/A - Integrated memory controller	90	Yes
Opteron 8YY (dual-core—for up to 8P servers)	2.4	64K data 64 K instruction	1024	None	N/A - Integrated memory controller	90	Yes

[2] The Microsoft press release dated October 19, 2004, states that "Microsoft software that is currently licensed on a per-processor model will continue to be licensed per processor, not per core, for hardware that contains dual-core and multicore processors." The press release is available at www.microsoft.com/presspass/press/2004/oct04/10-19multicorelicensingpr.mspx

As illustrated in Figure 1, the Opteron single-core and dual-core processors use the same basic architecture: The execution core(s) communicate with the memory controller and the HyperTransport links through the crossbar switch. The two execution cores include the processor core itself, the 64 KB L1 data and instruction cache, and the 1M L2 cache. The use of dual cores in the Opteron provides multi-threaded parallelism.

Figure 1. Schematic of a single-core and a dual-core Opteron die

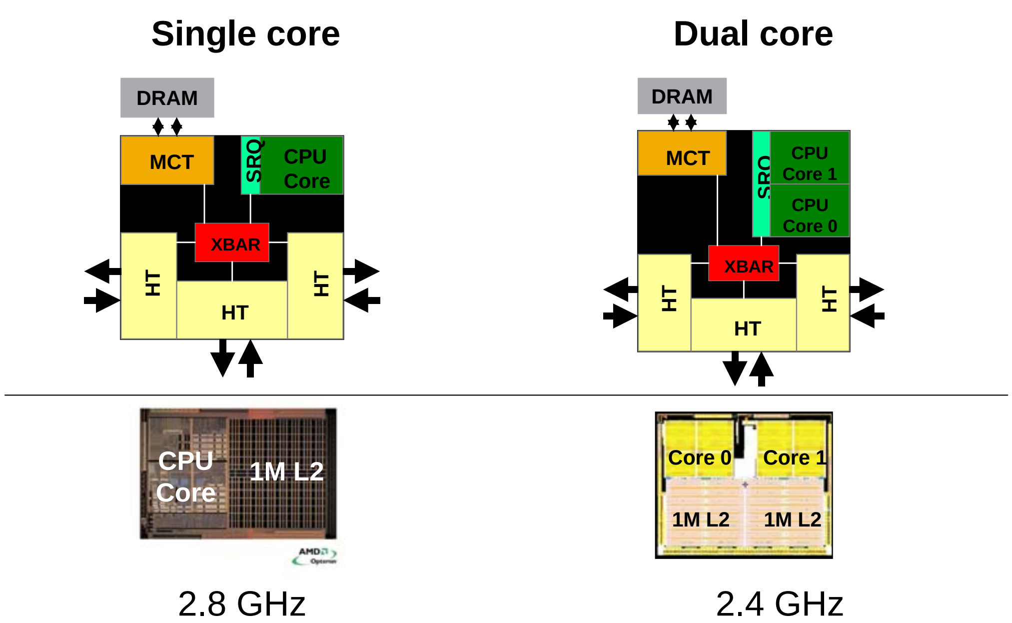

AMD processor naming conventions

A three-digit model number, XYY, identifies Opteron processor models. The X indicates the maximum scalability of the processor; that is, whether the processor is designed for single-, dual-, or multi-processor systems:[3]

- 100 series—single-processor servers. The HyperTransport links can only be used to connect to I/O, not to other processors.
- 200 series—up to two-processor servers. Only one HyperTransport link can be used to connect to another processor.
- 800 series—up to eight-processor servers. All three HyperTransport links can be used to connect to other processors.

The second two numbers, YY, indicate the relative performance within the processor series and indicate whether the processor is single-core or dual-core. Single-core processors are designated in factors of two, while dual-core processors are designated in factors of five. Thus, an AMD Opteron processor model 244 is a single-core DP processor with higher performance than an AMD Opteron processor model 242. An AMD Opteron processor model 875 is a dual-core MP processor with higher performance than the AMD Opteron processor model 865.

[3] For more information about naming of the AMD processors, see www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_9240,00.html.

Micro-architectural similarities

Both the Xeon family processors and the Opteron processors adhere to the x86 instruction set architecture to be compatible with the wealth of 32-bit software applications available. Therefore, at the programming level, both processors perform in the same way. In other words, at the software/hardware interface, the software interface of each processor remains the same with regard to the memory addressing size, the instruction sets, and the register designs for the x86 architecture.

32-bit operations

A 32-bit processor has general-purpose registers (GPRs) that are 32 bits wide and can operate on an integer data stream that is 32 bits wide. In addition, a 32-bit processor can hold 32 bits of memory address data in a single register, for a maximum of 4 GB of addressable memory.

The x86 architecture also supports physical addressing extensions (PAE), which extend the address space to allow addressing to 36 bits for a maximum of 64 GB of physical addressable memory. However, this requires the OS and applications to take advantage of the additional memory addressing.[4] The Xeon family processors and the Opteron processors support 32-bit addressing as well as the 36-bit PAE.

As shown in Table 3, the x86, 32-bit instruction set common to both the Xeon family processors and Opteron processors includes:

- Standard x86 instructions, which are general-purpose arithmetic functions
- Single Input Multiple Data (SIMD) Instructions, which allow one command to work simultaneously on multiple data items. This includes MMX, Streaming SIMD Extensions (SSE), SSE2, and SSE3 instructions.
- x87 floating point instructions

Table 3. 32-bit x86 instructions common to both Intel and AMD processors

Instruction name	Description	Register type	Size of registers	Number of registers
standard X86	Instructions for logical and arithmetic operations, address calculations, and holds memory pointers	GPR	32-bit	8
MMX	Multimedia instructions that allow the processor to do 64-bit SIMD operations	MMX	64-bit	8
x87	Instructions for floating point calculations	FP	80-bit*	8
SSE, SSE2, and SSE3	SSE improved upon the MMX instructions and allowed processors to do 128-bit SIMD floating-point operations. SSE2 added 64-bit parallel floating point numeric support. It also added new instructions to support 128-bit SIMD integer operations. SSE3 instructions include 13 instructions that accelerate performance of SSE technology, SSE2 technology, and x87-floating-point math capabilities.	XMM	128-bit	8

* According to the article "Introduction to 64-bit Computing and x86-64" by Jon "Hannibal" Stokes, the "x87 uses 80-bit registers to do double-precision floating point. The floats themselves are 64-bit, but the processor converts them to an internal, 80-bit format for increased precision when doing computations." [5]

[4] For more information, see "36-Bit Physical Addressing Using the PAE Paging Mechanism" in Chapter 3 of the IA-32 Intel Architecture Software Developer's Manual, Volume 3., available at www.intel.com/design/pentium4/manuals/index_new.htm

[5] Available at http://arstechnica.com/cpu/03q1/x86-64/x86-64-1.html

There is one minor difference between the 32-bit instruction sets: Opteron processors support the AMD 3Dnow!™ instructions; Intel processors do not. When Intel introduced the MMX instructions, they were not widely used, and AMD developed its own version of multimedia instructions with the 3DNow! instructions. The 3DNow! set added SIMD instructions to improve the vector-processing (floating point) requirements of graphic-intensive and multimedia applications. The 3DNow! instructions use the same registers as the MMX instructions.

64-bit operations

64-bit architectures have registers and arithmetic logic units that are capable of manipulating 64 bits of data during a single processing step. Because the registers store addresses to memory, this also means that a 64-bit architecture has a much larger amount of directly addressable memory than a 32-bit processor. Therefore, 64-bit architectures can provide performance advantages by their ability to use large amounts of memory (for example, in data mining operations) and by their ability to manipulate large amounts of numbers, such as with intensive floating-point calculations used in scientific and engineering modeling programs.

Instruction set and registers

As already noted, AMD was the first to develop the 64-bit extensions with their AMD64 instructions. Intel then delivered their EM64T instruction set that is broadly compatible with AMD64 for 64-bit computing. Both AMD64 and EM64T instructions can take advantage of the 64-bit wide registers in Opteron processors and the latest Xeon processors. These registers are used by the applications only when running the processors in 64-bit long mode. To support the AMD64 or EM64T instructions, the registers expand to include:

- Eight new 64-bit GPRs
- Extensions of the eight original, 32-bit GPRs to 64 bits
- Eight new 128-bit registers for SSE, SSE2, and SSE3 instructions

Figure 2. Difference between 32-bit and 64-bit registers

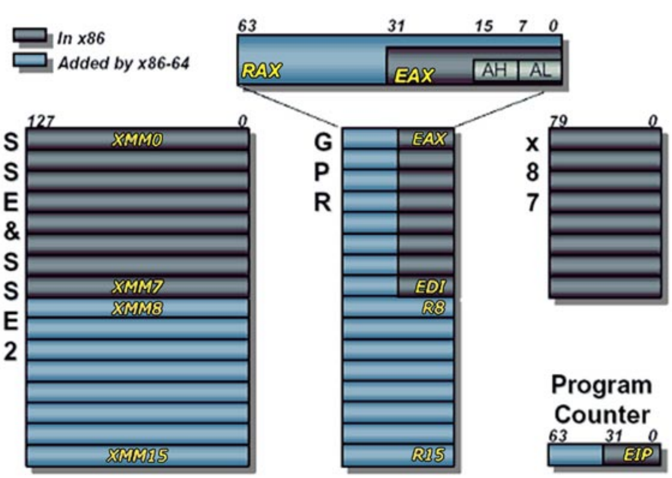

Source: "AMD Opteron – Part 1: Intro to Opteron/K8 Architecture" by Anand Lal Shimpi, April 23, 2003

Operating modes

The processors use three different operating modes: 64-bit long mode, 64-bit compatibility mode, and 32-bit legacy mode. The 64-bit long mode requires a 64-bit OS and an application recompiled to use the 64-bit registers. In other words, the full capabilities of the expanded register set are available only when both the OS and the application support 64 bits. The 64-bit compatibility mode requires a 64-bit OS, but can use a 32-bit application. The additional registers are available to the OS, but not to the 32-bit application, since it cannot make use of them. When running in legacy mode, the processor acts just like a 32-bit processor, and the extra registers are not available (Table 4).

Table 4. Operating modes for Xeon family and Opteron processors[6]

Mode	OS required	Application recompile required?	Register extensions available?	GPR width (bits)
64-bit long mode	64-bit OS	Yes	Yes	64
64-bit compatibility mode	64-bit OS	No	Yes – to OS No – to application	32
32-bit legacy mode	32-bit or 16-bit OS	No	No	32

Memory addressability

Although all the registers for the Xeon and Opteron registers are at least 64-bits wide, neither the Xeon family processors nor the Opteron processors use all 64 bits to address memory. Because this is such an extremely large number (2^{64} bits equals 16 exabytes), current applications cannot support this much memory in either local memory or page files.[7] In addition, the use of all 64 bits to address memory can cause the architecture to be more complex than it needs to be, wasting space and resources. Since the majority of memory accesses would not require such a large number to identify the address, having 64-bit memory addresses would result in lots of extra bits taken up by zeros. Therefore, when operating in 64-bit long mode, the Xeon family processors and the Opteron processors support up to 40 bits of physical memory (1 terabyte) and 48 bits of virtual memory (256 terabytes).

Micro-architectural differences

As just discussed, the Xeon family and Opteron processors have fundamentally similar designs at the programming level, meaning the instruction/register/addressing schemes are essentially the same at the software/hardware interface. Of course, there are many micro-architectural differences in the implementation of those processors (how the processor decodes instructions, how the branch prediction buffers work, how many branch predictions can be held in those buffers, etc.). However, it is important to understand two main differences between the processors:

- Overall design of the pipelines and how this relates to processor frequency. Xeon family processors optimize frequency at the expense of pipeline execution efficiency, while Opteron processors optimize the pipeline efficiency at the expense of frequency.
- Implementation of multi-threaded parallelism. Xeon family processors include Hyper-Threading technology, which allows different threads to operate in parallel, in addition to instruction-level parallelism. Opteron processors offer only instruction-level parallelism.

[6] From the documents titled "Intel Extended Memory 64 Technology Software Developer's Guide," Vol. 1, available at ftp://download.intel.com/technology/64bitextensions/30083402.pdf; and from "AMD64 Architecture Programmer's Manual, Vol. 1: Application Programming", available at www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/24592.pdf

[7] Appendix B defines some of the less familiar engineering prefixes such as exabytes.

Pipelines and clock frequency

The pipeline in a processor is analogous to an assembly line in a factory: to execute program code (or to build a widget), the work is split into multiple "stages," with each stage comprising a small part of the whole job. The idea, of course, is that splitting up the work into stages keeps the processor (or factory worker) busy at all times, allowing the processor to execute more code during a certain period of time.

It is important to understand the amount of activity that can occur within each stage of the pipeline. The processor must complete the operation for each stage within a single clock cycle. If the processor reduces the task size by splitting it into two or more smaller tasks, each stage can be shorter, but there will be more stages (Figure 3). Thus, each stage can be completed more quickly, allowing the processor to have a higher clock frequency. This does not necessarily mean that more work is being done in the pipeline; it just means that the clock frequency can be higher.

Figure 3. Decreasing the amount of work done in each stage allows the clock frequency to increase

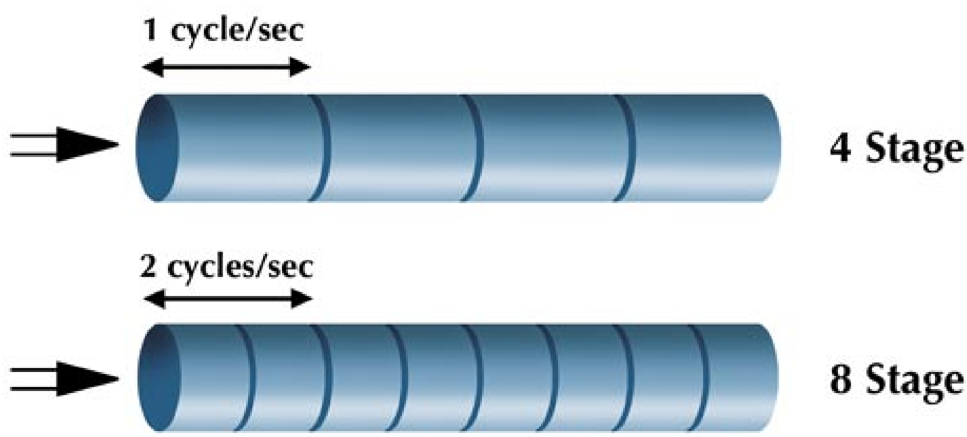

A basic structure for a computer pipeline consists of the following four steps, which are performed repeatedly to execute a program:

1. Fetch the next instruction from the address stored in the program counter.
2. Store that instruction in the instruction register and decode it, and increment the address in the program counter.
3. Execute the instruction currently in the instruction register.
4. Write the results of that instruction from the execution unit back into the destination register.

Typical processor architectures split the pipeline into segments that perform those basic steps: the "front end" of the microprocessor, the execution engine, and the retire unit, as shown in Figure 4. The front end fetches the instruction and decodes it into smaller instructions (commonly referred to as micro-ops). These decoded instructions are sent to one of the three types of execution units (integer, load/store, or floating point) to be executed. Finally, the instruction is retired and the result is written back to its destination register.

Figure 4. Basic 4-stage pipeline schematic

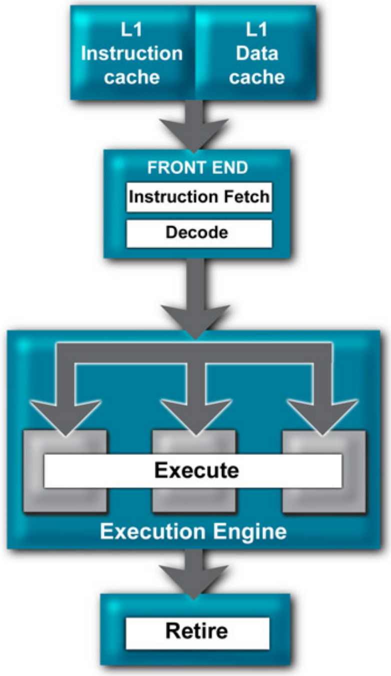

Processor stalls due to cache misses

Keeping the pipeline busy requires that the processor begin executing a second instruction before the first has traveled completely through the pipeline. However, suppose a program has an instruction that requires summing three numbers:

X = A + B + C

What happens if the processor already has A stored in a register, and B stored in a register, but needs to get C from memory? This causes a "bubble" or a stall in the pipeline, in which the processor cannot execute the instruction until it obtains the value for C from memory. This bubble must propagate all the way through the pipeline, forcing each stage that contains the bubble to sit idle, wasting execution resources during that clock cycle.

Clearly, the longer the pipeline, the more significant this problem becomes.

Processor stalls due to branch misprediction

Processor stalls often occur as a result of one instruction being dependent on another. If the program has a branch, such as an IF… THEN loop, the processor has two options. The processor either waits for the critical instruction to finish (stalling the pipeline) before deciding which program branch to take, or it predicts which branch the program will follow.

If the processor predicts the wrong code branch, it must flush the pipeline and start over again with the IF… THEN statement using the correct branch. The longer the pipeline, the higher the performance cost for branch mispredicts. For example, the longer the pipeline, the more the processor must execute speculative instructions which must be discarded when a mispredict occurs.

Xeon pipeline

As stated in documentation by Intel, the NetBurst architecture used in the Xeon processors provides a deep pipeline to enable industry-leading clock rates.[8] In other words, an important design criteria for Intel has been high processor clock rates. The most recent versions of the Xeon family processors use a *31-stage integer pipeline*, with two "fast" arithmetic logic units (ALUs), one "slow" ALU, two address generation units (AGUs), and two floating-point execution units (Figure 5). The long pipeline allows the Xeon processor to operate at frequencies up to approximately 3.8 GHz.

One of the ways that Intel has compensated for such a long pipeline is the 12KB instruction cache known as the Execution Trace Cache. Intel reports that this 12K cache has a hit rate similar to a conventional 8 – 16 KB instruction cache.[9] It stores translated and decoded instructions and puts them into traces, or "mini-programs." Whenever there is an L1 cache hit, the processor executes these traces without having to translate and decode the instructions for a cached trace, thus reducing the amount of work that the main pipeline must do. Furthermore, the Execution Trace Cache includes its own branch prediction algorithms so it can store translated micro-ops in speculative order.

Figure 5. Schematic of the Xeon processor microarchitecture

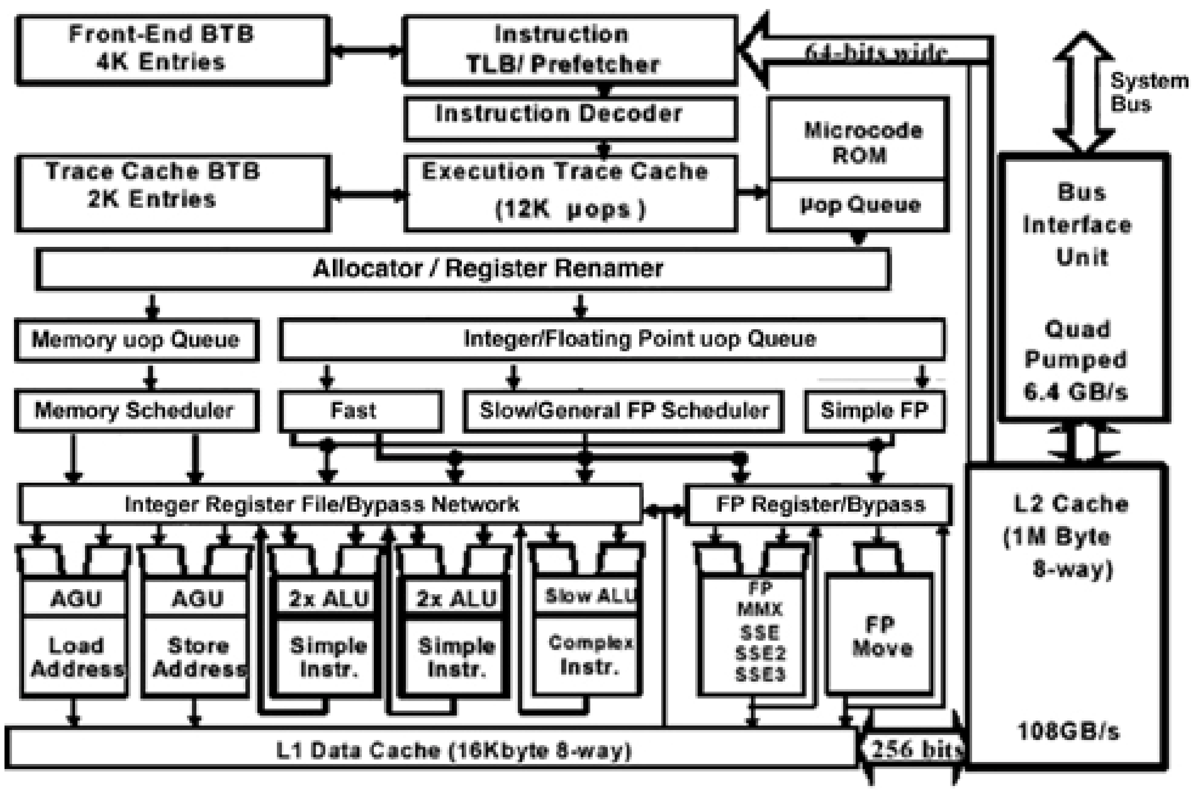

Source: "The MicroArchitecture of the Intel Pentium 4 Processor on 90nm Technology,"[10] Intel Technology Journal, Vol. 8, Issue 1, Feb. 2004.

[8] "IA-32 Intel Architecture Software Developer's Manual," Volume 1:Basic Architecture, pg. 2-7, available at http://developer.intel.com/design/pentium4/manuals/253665.htm

[9] Intel Technology Journal, Vol. 8, Issue 1, 2004, "MicroArchitecture of the Intel Pentium 4 Processor on 90 nm Technology."

[10] Available at http://www.intel.com/technology/itj/2004/volume08issue01/art01_microarchitecture/p01_abstract.htm

Opteron pipeline

Unlike the Xeon family processors, the Opteron processor has been optimized for a highly efficient pipeline at the expense of high clock speeds. According to AMD,[11] the pipeline's front end instruction fetch and decode logic has been optimized to pack multiple decoded, micro-op instructions together to execute them in parallel. The Opteron processor has a *12-stage integer pipeline*, much shorter than the pipeline for Xeon processors. The shorter pipeline requires a slower frequency. As of this writing, the Opteron processor operates at up to 2.8 GHz for single-core versions. However, the shorter pipeline reduces the risk of delays due to branch mispredicts and cache misses. The shorter pipeline also requires less extensive branch prediction algorithms and target buffers.

To make operations more parallel, the Opteron processor also has more execution units and decode units than Xeon processors. The Opteron processor includes three ALUs, three AGUs, and three floating-point execution units (Figure 6). Although it has more individual execution units than Xeon processors, the maximum effective throughput of the Opteron execution units is the same as in a Xeon processor—three integer operations per cycle.

Figure 6. AMD Opteron architectural block diagram

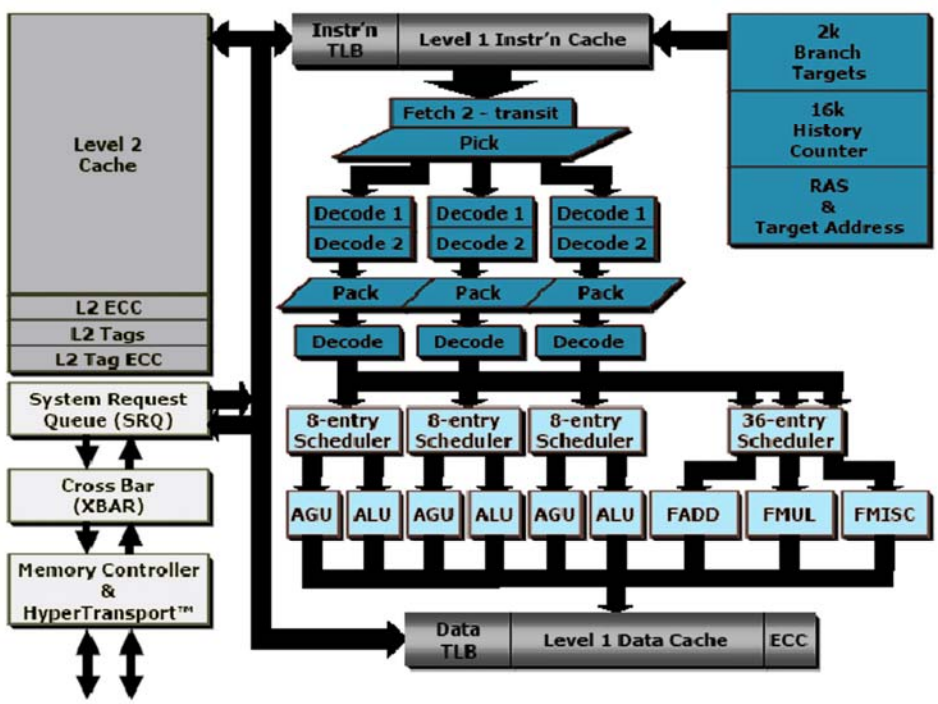

Source: AMD white paper "AMD Eighth-generation Processor Architecture," October 16, 2001

[11] "AMD Eighth-generation Processor Architecture," available at www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/Hammer_architecture_WP_2.pdf

Pipeline and frequency comparison

In contrast to the Xeon processors which have high frequencies and long pipelines, several processors designed to achieve high performance—such as the Alpha EV68, HP-PA RISC, and the Itanium processors—tend to be designed with shorter pipelines and to operate at relatively lower frequencies. For example, the Alpha EV68/21264 processor uses a 7-stage pipeline and operates at approximately 1.25 GHz; and the Itanium 2 processor uses an 8-stage pipeline and operates at approximately 1.6 GHz. This does not necessarily mean that Xeon processors have lower performance because they have a longer pipeline. However, it does mean that the Xeon processors must compensate for the long pipeline by techniques such as more efficient branch prediction algorithms and larger look-aside buffers. It also means that Xeon processors may be best suited for applications where high clock speed is especially important, such as applications using a linear programming style.

Hyper-Threading in the Xeon family processors

The Intel NetBurst architecture incorporates Hyper-Threading as another method to keep its pipeline full. Normally, a processor executes instructions from a single instruction stream. Whenever a switch occurs between threads, or streams of instructions, the processor must pay an overhead penalty for the context switch. Because modern processors are superscalar (have parallel execution units and out-of-order execution), the processors re-arrange the instructions and execute them out-of-order to perform some instructions in parallel (instruction-level parallelism). Both the Xeon and the Opteron processors use instruction-level parallelism to execute up to three instructions per clock.

In addition, the Xeon family processors use Hyper-Threading technology to execute two separate threads in parallel (multi-threaded parallelism). The Xeon family processors replicate certain components of the processor such as the instruction pointer, register allocation tables, and other architectural registers. Intel refers to these components as the “architectural state.” The architectural states, or logical processors, schedule instructions for the processor execution resources, which are shared between the two logical processors. When the processor identifies a bubble in one thread stream, it shifts the second thread stream into those execution resources. This allows the processor to execute incoming micro-op instructions from different threads in a time-multiplexed manner to keep execution resources as busy as possible (Figure 7).

The benefits of Hyper-Threading are directly related to the deep execution pipeline of the Xeon NetBurst micro-architecture. Because the effects of pipeline stalls (bubbles) are exaggerated in a deep execution pipeline, the Xeon architecture implemented Hyper-Threading technology to provide high levels of multithreading performance by keeping execution units as busy as possible. A shorter, more efficient micro-architecture such as that implemented in the AMD Opteron processor would not benefit from Hyper-Threading as much as the hyper-pipelined Xeon micro-architecture, nor does it require this level of complexity to provide high levels of multithreaded performance.

Figure 7. Comparison of single-core processors with and without Hyper-Threading technology In the Hyper-Threading example on the right, the light gray indicates instructions from thread 1, and the dark gray indicates instructions from thread 2. The white squares are idle execution resources.

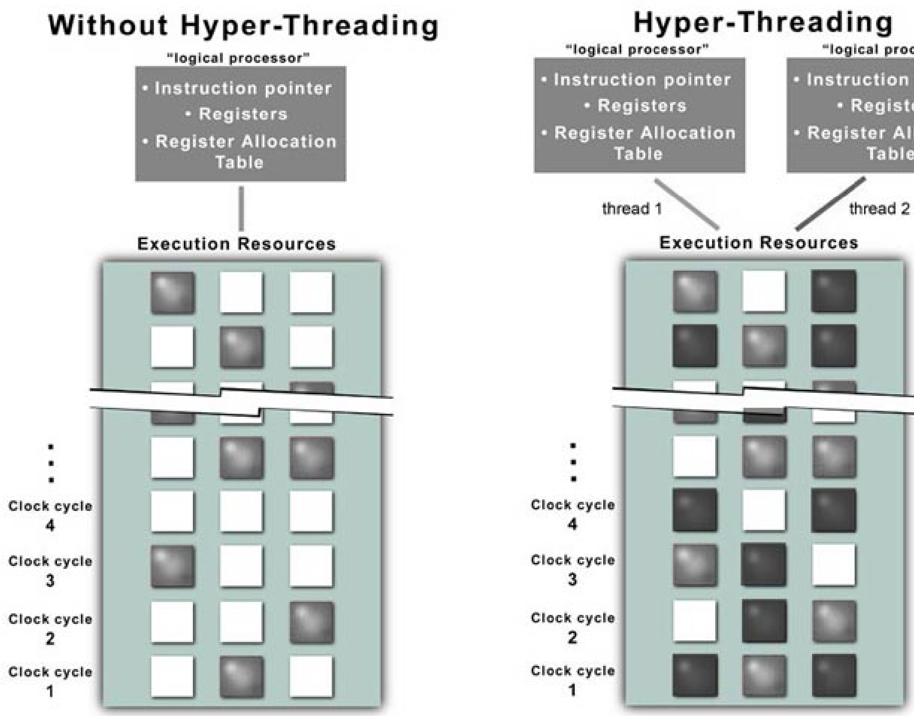

Hyper-Threading can improve system performance for applications and operating systems that can take advantage of multi-threading. Intel reports maximum gains of up to 30 percent,[12] but this is highly dependent on the application. Average performance improvements are more likely to be approximately 5 to10 percent,[13] and depend on how well the OS and applications take advantage of multiple threads. In other words, applications and operating systems that are not multi-threaded may not gain any performance benefits from Hyper-Threading.

Dual-core Intel processors designed for the server market are expected to include Hyper-Threading, so that a single-processor, dual-core system will have the potential to execute four threads seemingly in parallel. Because there are two separate execution cores, two instruction threads can be executed simultaneously; and two other threads can use the Hyper-Threading resources to access idle execution resources when a bubble occurs. Again, actual performance gains will depend on how well the OS and applications can use multiple threads.

[12] Source:"Hyper-Threading Technology for Servers," available at http://www.intel.com/business/bss/products/hyperthreading/server/index.htm

[13] The Anandtech website reports 3 to 5 percent gains in database performance testing; "AMD Opteron vs. Intel Xeon: Database Performance Shootout," March 2, 2004, http://www.anandtech.com/IT/showdoc.aspx?i=1982

System architecture differences

Server performance depends upon both the processor micro-architecture—which has already been discussed—and the processor macro-architecture, or how the processor interfaces with the memory and I/O subsystems of the server. At the system level, there are two main differences between the Xeon architecture and the Opteron architecture:

- The Xeon family processors use a parallel, shared, front-side bus to access a separate chipset that connects to the memory and I/O subsystems.
- Opteron processors include their own integrated memory controller that provides a direct connection to memory. The processors connect to the I/O subsystem using point-to-point HyperTransport links that can offer significantly higher bandwidth and lower latencies than a shared front-side bus.

This section discusses only in general terms how system architecture might vary based on processor choice. For specific information about ProLiant platforms, see the HP website at www.hp.com/go/proliant.

Intel architecture

Whether a system uses one or eight Xeon family processors, all communication requests between those processors and the memory or I/O subsystems must travel from the processor, across one or two front-side buses, to a north bridge chip, and then to memory or to a south bridge chip, which connects to I/O. Therefore, the bandwidth, scalability, and latencies of the memory and I/O subsystems are highly dependent upon the chipset design as well as the attributes of the front-side bus.

Bandwidth

The Intel front-side bus (FSB) is a parallel, 64-bit, multi-drop technology that shares bandwidth equally among the processors on that bus. ProLiant platforms available as of October 2005 use Xeon family processor generations that have effective front-side bus speeds varying between 400 and 800 MHz. This results in a maximum bandwidth between 3.2 and 6.4 GB/s along the FSB.

Figure 8 shows a representative example of a 2P server architecture using the Xeon processor: All memory traffic goes through the north bridge, acting as the sole memory controller; and all I/O traffic continues through a hub or a south bridge to the I/O devices.

Figure 8. Example of a 2P server architecture using Xeon processors

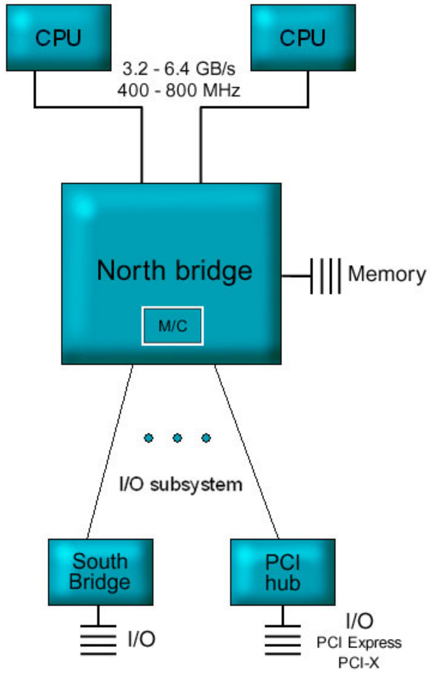

Recent 4P ProLiant platforms using the Xeon MP processor family have incorporated an Intel chipset that uses two front side buses rather than a single front side bus and multiple memory controllers rather than a single north bridge (Figure 9). The multiple memory controllers allow concurrent memory accesses, which improves memory performance. However, the front side buses can still be limiting factors because the chipset requires that processors from one bus must snoop the other bus whenever a memory request occurs.

Figure 9. Example of a 4P- architecture using Xeon MP processors and dual front-side buses

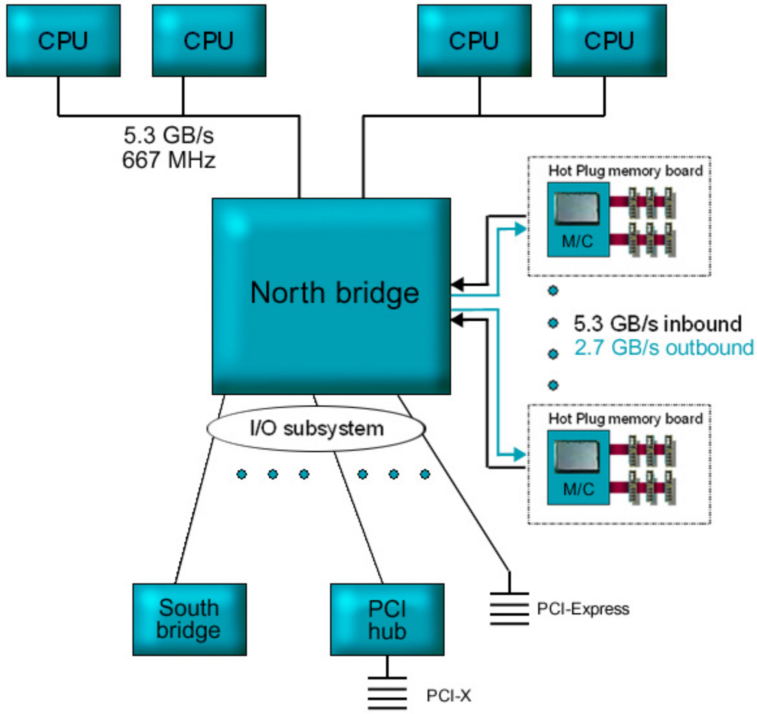

Latency

The Xeon family processors are designed for use in a symmetric multi-processing environment in which every processor has equal access to the memory. Therefore, memory latency—the time it takes for a processor to request data from memory—is uniform across all processors.

However, Xeon processors tend to have larger memory latencies than Opteron processors due to the speed of the memory controller. The external chipset used with Xeon and Xeon MP processors operates in the range of hundreds of megahertz, while the internal memory controller used in the Opteron processor operates in the gigahertz range.

Scalability

As customers add more processors, scalability is determined by the ability of the processors to share the limited FSB bandwidth, especially with applications that are memory intensive. The shared nature of the FSB results in frequent arbitration and thus higher latencies, which may limit how well the Xeon family processors can take advantage of those memory resources. Additionally, the FSB speeds available with Xeon MP processors are typically lower than those with Xeon processors, further constraining the available FSB bandwidth. Thus, performance does not scale up linearly as the number of processors increases. To improve performance scalability in multiprocessor (4P) platforms, Intel provides Xeon MP processors with large caches to store more data within the processor. This reduces the need for the processor to use the FSB, thereby improving performance.

AMD architecture

As already mentioned, the Opteron processor integrates the memory controller into the processor; therefore, memory is attached locally to each processor. In an SMP system, other processors can access that memory by means of a crossbar switch inside the processor and fast, point-to-point HyperTransport interconnects between processors (Figure 10). Each Opteron processor contains three HyperTransport links that can connect to other processors or to the I/O subsystem, depending on whether it is a 100-series, 200-series, or 800-series processor (as described in the section titled "AMD processor naming conventions").

Figure 10. Representative example of a 4P architecture using Opteron processors

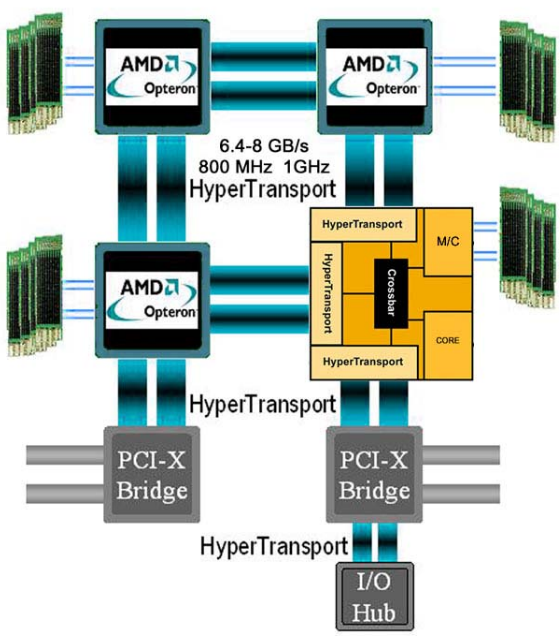

HyperTransport is a parallel point-to-point interconnect that replaces parallel front side bus technology. It uses a double-data rate to transfer 16 bits of data on both the rising and falling edges of the clock signals, resulting in an effective 32 bits of data per clock cycle on the 16-bit link. Opteron processors use a HyperTransport clock speed of 800 MHz to 1 GHz. A processor that uses 1-GHz HyperTransport links provides an effective transfer rate of up to 4 GB/s in each direction. Since transfers can occur in both directions simultaneously, an aggregate transfer rate of 8 GB/s can be achieved in each16-bit HyperTransport I/O link. Compared to a shared or bi-directional bus, a point-to-point interconnect has the advantage of no overhead for bus arbitration and easier maintenance of signal integrity.

Because of the integrated memory controller and point-to-point HyperTransport links, systems using Opteron processors tend to provide more bandwidth, less latency, and better scalability than systems using Xeon family processors.

Bandwidth

Because the memory controller is integrated into the processor, all the data paths inside the controller run at core frequencies (1.8 to 2. 8 GHz, depending on the processor). Each on-chip memory controller directly accesses its associated DIMMs at the DIMM speed—up to 400 MHz when using

PC3200 DDR memory. For PC3200 memory, this gives a bandwidth of up to 3.2 GB/s per channel, or a total of 6.4 GB/s for both channels to the processor. Therefore, systems using Opteron could have up to 100 percent greater bandwidth than systems using a Xeon family processor, depending on the front-side bus speed of the Xeon family processor.

The bandwidth of the I/O subsystem depends upon platform design and the chipset used. The AMD 8000 chipset normally used with the Opteron processors includes HyperTransport tunnels that can provide 4.8 GB/s to 8 GB/s of bandwidth, depending on whether 8-bit or 16-bit links are used. Even an 8-bit HyperTransport link provides adequate headroom for future expansion and high I/O throughput, as shown by the comparisons in Table 5.

Table 5. Bandwidth comparison between HyperTransport and existing I/O protocols

Protocol	Bandwidth	HyperTransport (8-bit, 4.8 GB/s) is faster by:	HyperTransport (16-bit, 8 GB/s) is faster by:
legacy PCI (32 bit, 33 MHz)	133 MB/s	36X	60X
USB 2.0	480 MB/s	10X	16.6X
PCI-X (64 bit, 133 MHz)	1064 MB/s	4.5X	7.5X
PCI Express (x8 link)	4 GB/s	1.2X	2X
InfiniBand 4X link	10 Gb/s (1.25 GB/s)	3.8X	6.4X
10 Gb Ethernet	10 Gb/s (1.25 GB/s)	3.8X	6.4X

Latency

Some customers may be concerned that the AMD architecture, which splits the memory into a processor's local memory and memory that is remote to the processor, will result in large latencies for remote memory accesses. However, the latency difference between local and remote accesses is actually very small because the memory controller is integrated into the processor and operates at the core speed of the processor, and because of the fast interconnect between processors. The difference between local and remote memory accesses is so small that it is sometimes referred to as "sufficiently uniform" memory. Although the memory subsystem design is not uniform, the speed of the interconnects allows the latencies to appear uniform.

However, because memory is local to the processors and the HyperTransport links can transfer data from only one request at a time, latencies can increase as memory requests increase, especially if the requests must "hop" across multiple HyperTransport links. In some operating systems, the OS has the ability to take advantage of local and remote memory so that memory latencies can be further reduced. HP ProLiant platforms such as the ProLiant DL585 server have a setting in the ROM-Based Setup Utility to exploit local memory on the processor rather than going through the crossbar switch in the processor to remote memory.

In general, HP measurements show that the memory latency for an Opteron-based system is approximately 10 to 40 percent less than that of a Xeon-based system. Another way to look at latency is to compare the difference in number of clock cycles that it takes processors to execute a memory request. Figure 11 shows that an Opteron processor uses significantly fewer clock cycles to achieve a memory transaction than an Intel processor. Although Figure 11 was developed using previous generations of processors, the relationship is still valid: regardless of the clock speed, every extra cycle spent on a memory access is a clock cycle that is not being used for other useful work.

Figure 11. Number of clock cycles spent waiting for memory accesses

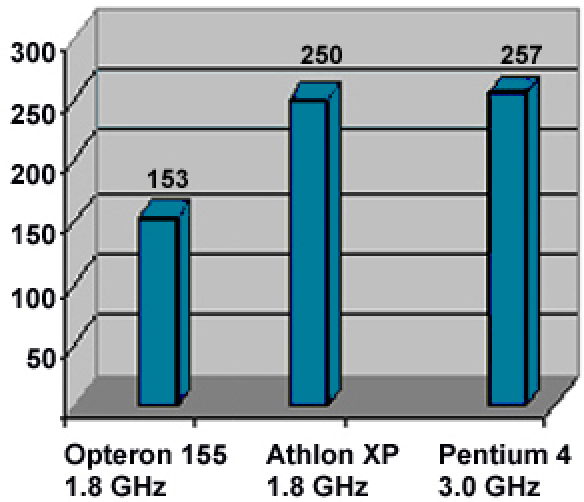

Source: "AMD Opteron Coverage – Intro to Opteron/K8 Architecture," April 23, 2003[14]

Scalability

Because customers can add memory with each processor in an Opteron-based system, total memory can scale linearly with the number of processors. For example, a ProLiant DL585, 4P Opteron system can have up to 128 GB of memory. For applications that can use this amount of memory, an Opteron-based system could provide significant performance advantages.

In addition, because each Opteron processor contains its own memory controller and HyperTransport links, the available memory bandwidth and processor-to-processor bandwidth scales linearly with the number of processors. The integrated memory controllers also allow multiple memory requests to be made in parallel. The parallelism increases the effective bandwidth to memory and decreases the average memory latency.

Power considerations

Power consumption is a growing challenge for large datacenters and customers using densely packed racks. The processors clearly require a large percentage of the power budget for a server. Both AMD and Intel are aware of this issue and are working to develop high-performing processors that use less power.

[14] Available at www.anandtech.com/cpuchipsets/showdoc.aspx?i=1815

Power consumption is directly related to the operating frequency of the processor: as the frequency increases, the power also increases. As of this writing, Opteron processors generally run at lower operating frequencies than Xeon processors. Therefore, Opteron processors will typically require less power and run cooler than Xeon family processors. Other factors that affect power consumption include cache size, transistor count, and silicon process technology.

While the Intel roadmap calls for lower-power processors, the current trend for high-performance Xeon and Xeon MP family processors shows continued increases in power consumption until the second half of 2006. For example, the previous generation of Xeon family processors (code-named Nocona and Gallatin) operated in the 100 W range. The top-end Xeon family processors now operate at up to approximately 130 W. By contrast, both dual-core and single-core Opteron processors operate at a maximum of approximately 95 W.

Comparing 32-bit performance

The difficulty in comparing architectures is that there are no absolutes about which processor will have better performance in a given system or application. For example, processor performance is a function of core processor frequency, pipeline design, branch prediction capabilities, cache sizes, and the interface to memory and I/O.

In many cases, servers using the Opteron processor can provide performance improvements when compared to servers using the Xeon family of processors. This is especially true for applications in which frequent memory accesses take advantage of Opteron's integrated memory controller. The performance differences are found most frequently in 4P systems because of the cumulative effects of multiple, integrated memory controllers and point-to-point HyperTransport links. Performance in 2P systems is much more dependent on the application used.

In addition, the advent of dual-core AMD Opteron processors has enabled ProLiant servers with dual-core processors to achieve significantly improved performance while maintaining the same server infrastructure and power budget. For example, a SPECWeb99 benchmark showed that performance of the 2P DL385 server increased by almost 50 percent after switching from single core to dual-core Opteron processors.[15] It is expected that when Intel's next-generation dual-core processors are released, server performance in Xeon-based systems will also improve significantly.

For the most accurate comparisons, customers should perform their own testing to compare Xeon- or Opteron-based systems. Customers can also evaluate the latest published benchmark results for each architecture. However, benchmarks only provide snapshot pictures that indicate the relative level of performance a customer can expect—they are not meant to duplicate customer situations exactly. For the most recent benchmark information for ProLiant servers, see the HP website. [16]

Summary

HP provides customers with a wide choice of platforms using either Intel Xeon processors or AMD Opteron processors. Both are 32-bit processors that comply with the x86 architecture. For customers that want or need to move to 64-bit applications and operating systems, both the Xeon and Opteron processors enable a smooth transition with 64-bit extensions while remaining compatible with existing 32-bit applications. This gives the Xeon and Opteron processors a distinct advantage for customers who want to access the extra level of memory addressability available with these processors in 64-bit mode, without the cost and infrastructure changes required to move to a 64-bit Itanium system.

The Xeon and Opteron processor architectures differ in their internal design details and platform implementations. The Xeon pipeline structure is optimized for high clock speeds, while the Opteron

[15] Refer to the Performance Brief titled "HP ProLiant DL385 fastest 2P server on the SPECweb99_SSL Benchmark," available at ftp://ftp.compaq.com/pub/products/servers/benchmarks/dl385-specweb99_ssl.pdf
[16] ProLiant server benchmark information is available at http://h18004.www1.hp.com/products/servers/benchmarks/index.html

pipeline is optimized for parallel execution. In the platform architecture, Xeon processors use the familiar front-side bus technology with a separate chipset for the memory controller(s). The Opteron processors use an integrated memory controller that runs at core processor speed. The Opteron direct memory interface enables large bandwidths and low latencies. The Opteron processor uses a direct HyperTransport point-to-point interconnect for the I/O subsystem, which provides high bandwidths. AMD's use of the integrated memory controller and the point-to-point I/O connections appear to greatly improve performance, allowing Opteron-based systems to outperform Xeon-based systems in many cases, especially in 4P systems. Recent testing shows that the 4P ProLiant DL585 server offers outstanding performance in typical benchmarks in transaction processing, messaging, and SAP database configurations.

Appendix A. Engineering prefixes

Table A1. Common engineering prefixes

	Abbreviation	Exponential form	Number of bytes	Relationship to next lowest prefix
Gigabyte	G/GB	2^{30} bytes	1,073,741,824 bytes	1024 Megabytes
Terabyte	T/TB	2^{40} bytes	1,099,511,627,776 bytes	1024 Gigabytes
Petabyte	P/PB	2^{50} bytes	1,125,899,906,842,624 bytes	1024 Terabytes
Exabyte	E/EB	2^{60} bytes	1,152,921,504,606,846,976 bytes	1024 Petabytes

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
AMD website	www.amd.com/
Opteron technical documents	www.amd.com/us-en/Processors/DevelopWithAMD/0,,30_2252_739_9003,00.html
AMD64 technical documents	www.amd.com/us-en/Processors/DevelopWithAMD/0,,30_2252_739_7044,00.html
Anandtech website Contains information about general processor architecture, Intel architecture, and AMD architecture	www.anandtech.com/
Ars Technica website: Contains information about general processor architecture, Intel architecture, and AMD architecture	www.arstechnica.com
HP website	
ProLiant servers	www.hp.com/go/proliant
Industry-Standard Server Technology Papers	www.hp.com/servers/technology
AMD processor roadmap for HP ProLiant servers, Technology Brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00428708/c00428708.pdf
Intel processor roadmap for industry-standard servers, Technology Brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00164255/c00164255.pdf
Intel Hyper-Threading Technology, Technology Brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00164255/c00164255.pdf
Intel website	www.intel.com/
Hyper -Threading Information	www.intel.com/technology/hyperthread

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com

TC051008TB, 10/2005